Supporting Material for 802.1AS D6.1 Ballot Comments 35, 36, 37, and 80

Geoffrey M. Garner
SAMSUNG Electronics (Consultant)

IEEE 802.1 AVB TG
2009.09.08

gmgarner@comcast.net

Comment 35 - 1

- Comment (subclause B.3, page 220, line 37)
  - B.3 states that the requirements of this standard and of standards referenced for each medium enable a time synchronization requirement of 1 us (+/- 500 ns) over 7 hops to be met. It would be useful to make a similar statement for jitter/wander requirements.
- Suggested remedy
  - Add a subclause to annex B entitiled "End-to-end jitter and wander performance". Suggested text, and an accompanying figure and table, are on the following slides.

Comment 35 - 2

B.4 End-to-end jitter and wander performance

The requirements of this standard and standards referenced by this standard ensure that the synchronized time at a time-aware system that is separated from the grandmaster by six or fewer time-aware systems (i.e., seven or fewer hops) will, when filtered by a reference endpoint filters with rolloff of 20 dB/decade, gain peaking that does not exceed 0.1 dB, and bandwidth that does not exceed the value given in each entry of Table , have wander maximum time interval error (MTIE) that does not exceed the maximum time interval error (MTIE) for that entry of Table B-4, and jitter that does not exceed the peak-to-peak jitter of Table B-4.

Note – For example, the endpoint filter can be of the form

$$y_k = a_1 y_{k-1} + a_2 y_{k-2} + \cdots + a_n y_{k-n} + b_0 x_k + b_1 x_{k-1} + \cdots + b_n x_{k-n} \quad , \qquad \text{(B-4)}$$

where the x_k are the unfiltered synchronized time values, the y_k are the filtered synchronized time values, and the a_k and b_k are filter coefficients. The a_k and b_k are chosen such that the filter has desired bandwidth and gain peaking that does not exceed 0.1 dB. Eq. (B-4) is a general infinite impulse response (IIR) digital filter. Simplified forms, e.g., a second order IIR filter obtained by setting n = 2, or a finite impulse response (FIR) filter obtained by setting the a_k to zero are possible.

Comment 35 - 3

Table B-4 – Maximum endpoint filter bandwidths needed to meet respective MTIE masks

Endpoint filter maximum bandwidth (Hz)	Corresponding wander MTIE mask of Figure B-4 that is not exceeded	Corresponding jitter high-pass measurement filter (Hz)	Corresponding peak-to-peak jitter\ (ns)
10	Mask 2	8000	10.2
1	Mask 1	200	11.1
0.01	Mask 3	–	–

Comment 35 - 4

Table B – 5. Breakpoints for mask 1

Observation Interval S (s)	MTIE] (ns)
$0.05 \le S \le 0.0637$	6954.8S
$0.0637 \le S \le 0.3183$	443
$0.3183 \le S \le 10000$	50000S

Table B – 6. Breakpoints for mask 2

Observation Interval S (s)	MTIE] (ns)
$0.05 \le S \le 0.4069$	407
$0.4069 \le S \le 10000$	1000S

Comment 35 - 5

Table B – 7. Breakpoints for mask 3

Observation Interval S (s)	MTIE] (ns)
$6.67 \times 10^{-4} \leq S \leq 4.0$	50S
$4.0 \leq S \leq 10000$	200

Comment 35 - 6

Figure B – 4. Wander MTIE masks

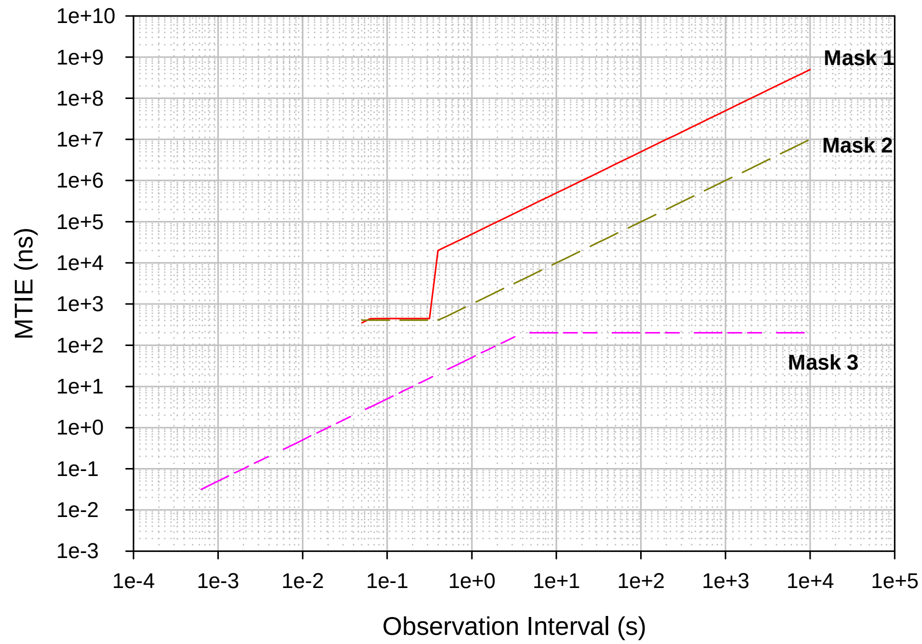

Comment 35 - 6

- Additional notes (not part of proposed text, but could be used if desired)
  - The first entry in Table B – 4, with 10 Hz endpoint filter, corresponds to the requirements for professional audio
  - The second entry in Table B – 4, with 1 Hz endpoint filter, corresponds to the requirements for consumer audio (the filter BW is narrower than for professional audio because consumer audio jitter tolerance requirement is lower)
  - The third entry in Table B – 4, with 0.01 Hz endpoint filter, corresponds to BTS (e.g., CDMA2000, WCDMA, femtocell)
  - Uncompressed video requirements are not included, pending outcome of liaison to SMPTE
  - If desired, we could indicate explicitly what applications these masks and jitter requirements correspond to, but we would then have to reference the relevant documents in the bibliography

Comment 35 - 7

- Additional notes continued (not part of proposed text, but could be used if desired)
  - MTIE masks 1 and 2 have been shown only for the wander region (i.e., minimum observation interval of 0.05 s, which corresponds to sampling frequency of 10 Hz)
    - Jitter requirements are shown in terms of peak-to-peak jitter with respective high-pass measurement filter
    - These jitter requirements are equivalent to the MTIE masks for shorter observation intervals used previously (those masks were originally derived from these jitter requirements (see http://www.ieee802.org/3/re_study/public/200505/garner_3_0505.pdf

Comment 36 - 1

- Comment (subclause 11.1.2, page 117, line 14)
  - Discussion in the AVB TG has indicated that the use of propagation time averaging will improve time synchronization performance and may reduce MTIE due to a frequency step at the grandmaster or grandmaster instability.
- Suggested remedy
  - Add a description of propagation time averaging to the general discussion of propagation delay measurement in 11.1.2.
  - Suggested text is shown on the next slide; it would be inserted just before the final paragraph of 11.1.2

Comment 36 - 3

to the responder, otherwise there will be an error equal to $0.5y(t_3 - t_2)$, where y is the frequency offset of the initiator relative to the responder. Finally, if the propagation delay measurement is desired relative to the grandmaster timebase, each term must be multiplied by the rate ratio of the grandmaster relative to the time base that term is expressed in.

There can also be an error in measured propagation delay due to time measurement granularity (see B.1.2). For example, if the time measurement granularity is 40 ns (as specified in B.1.2), the timestamps t_1, t_2, t_3, and/or $t4$ can undergo 40 ns step changes. When this occurs, the measured propagation delay, D, will change by 20 ns (or by a multiple of 20 ns if more than one of the timestamps has undergone a 40 ns step change). The actual propagation delay has not changed by 20 ns; the effect is due to time measurement granularity. The effect can be reduced, and the accuracy improved, by averaging successive measured propagation delay values. For example, an exponential averging filter can be used, i.e.,

$$D_{avg,k} = aD_{avg,k-1} + (1-a)D_{k-1}, \qquad (11\text{-}2)$$

where D_k is the k^{th} propagation delay measurement, $D_{avg,k}$ is the k^{th} computed average propagation delay, and k is an index for the propagation delay measurements (i.e., peer delay message exchange). The quantity a is the exponential weighting factor; it can be set so that the weight of a past propagation delay measurement is $1/e$ after M measurements, i.e.,

$$a = e^{-\frac{1}{M}}. \qquad (11\text{-}3)$$

The above averager must be initialized. One method is to use a simple average (i.e., the sum of the number of samples divided by the number of samples) of the measurements made up to the current measurement until a window of M measurements has been accumulated. In this case, Eq. (11-2) is used only for $k > M$. For $k \le M$, the averaged propagation delay is given by

$$D_{avg,k} = \frac{(k-1)D_{avg,k-1} + D_{k-1}}{k}. \qquad (11\text{-}4)$$

The rate ratio of the responder relative to the initiator is the quantity neighborRateRatio (see 10.2.7.5). It is computed by the function computePdelayRateRatio() (see 11.2.15.2.3) of the MDPdelayReq state machine (see 11.2.15) using successive values of t_3 and t_4. As indicated in the description of computePdelayRateRatio(), any scheme that uses this information is acceptable as long as the performance requirements of B.2.3 are met. One example scheme is given in Note 1 of 11.2.15.2.3.

Comment 37 - 1

- Comment (subclause 11.2.15.2.4, page 132, line 28)
  - Discussion in the AVB TG has indicated that the use of propagation time averaging will improve time synchronization performance and may reduce MTIE due to a frequency step at the grandmaster or grandmaster instability..
- Suggested remedy
  - Add a description of propagation time averaging to the function computePropTime().
  - Suggested text is shown on the next slide; it would be inserted just before Note 1

11.2.15.2.4 computePropTime(): computes the mean propagation delay on the link attached to this MD entity, D, and returns this value. D is given by

$$D = \frac{r \cdot (t_4 - t_1) - (t_3 - t_2)}{2}, \qquad (11\text{-}6)$$

where:

t_4 = <pdelayRespEventIngressTimestamp> (see 11.3.2.1) for the Pdelay_Resp message received in response to the Pdelay_Req message sent by the MD entity, expressed in ns; the <pdelayRespEventIngressTimestamp> is equal to the time stamp value measured relative to the time stamp measurement plane, minus any ingressLatency (see 8.4.3);

t_1 = <pdelayReqEventEgressTimestamp> (see 11.3.2.1) for the Pdelay_Req message sent by the P2PPort entity, expressed in ns,

t_2 = sum of (i) the ns field of the requestReceiptTimestamp, (ii) the seconds field of the requestReceiptTimestamp multiplied by 10^9, and (iii) the correction field divided by 2^{16} (i.e., the correction field is expressed in ns plus fractional ns), of the Pdelay_Resp message received in response to the Pdelay_Req message sent by the MD entity,

t_3 = sum of (i) the ns field of the responseOriginTimestamp, (ii) the seconds field of the responseOriginTimestamp multiplied by 10^9, and (iii) the correction field divided by 2^{16} (i.e., the correction field is expressed in ns plus fractional ns), of the Pdelay_Resp_Follow_Up message received in response to the Pdelay_Req message sent by the MD entity, and

r = current value of neighborRateRatio for this MD entity (see 10.2.7.5).

Propagation delay averaging may be performed, as described in 11.1.2 by Eqs. (11-2), (11-3), and (11-4). In this case, the successive values of propagation delay computed using Eq. (11-6) are input to either Eq. (11-2) or Eq. (11-4), and the computed average propagation delay is returned by this function.

NOTE 1 - Eq. (11-6) defines D as the mean propagation delay relative to the time base of the time-aware system at the other end of the attached link. It is divided by neighborRateRatio (see 10.2.7.5) to convert it to the time base of the current time-aware system when adding to <syncEventIngressTimestamp> to compute rxTime (see 11.2.13.2.1 (f)).

Comment 80 - 1

❑Comment

- The derivation of the 1 ns peak-to-peak jitter requirement was based on a noise generation model used in previous simulations. The model contained a combination of white phase modulation (WPM), flicker phase modulation (FPM), and flicker frequency modulation components (FFM). The FFM component is dominant at low frequencies, i.e., less than approximately 1 kHz. The WPM conponent was dominant at high frequencies, i.e. greater than about 10 kHz. The FPM component was mainly appreciable at intermediate frequencies, i.e., 10 - 100 kHz. The WPM level at high frequencies was considerably below the FFM level around 1 kHz, i.e., -40 dBc for the former vs -120 dBc for the latter. Unfortunately, the jitter limit was derived from the WPM component, but assumed to be measured through a 10 Hz high-pass filter. The resulting jitter limit obtained from the -120 dBc level was considerably less than 1 ns rms; it was rounded to 1 ns pp because, in addition to this jitter, there is the much larger 40 ns possible phase error due to phase measurement granularity. In any case, the derivation of this 1 ns, measured through a 10 Hz high-pass filter, is incorrect and should be revisisted.

Comment 80 - 2

❑Comment – cont.

- It also has been pointed out to the editor that the 10 Hz high-pass filter may not be practical to test with, using existing test equipment. A value of 12 kHz has been suggested to the editor (all this discussion has been offline). Note that 12 kHz as a high-pass jitter measurement filter is used to measure SONET (i.e., SDH Option 2) jitter generation; this is a long-standing measurement in SONET/SDH Option 2. The filter is specified in both ANSI T1.105.03 and ITU-T Rec. G.783, for OC-1, OC-3 (STM-1), OC-12 (STM-4), and OC-48 (STM-16). However, the filter is NOT used in SDH Option 1 (Option 1 is for networks optimized for the 2048 kbit/s hierarchy; Option 2 for networks optimized for the 1544 kbit/s hierarchy). Also, in SONET and SDH Option 2 it is NOT used for network interface jitter measurements, nor does a 12 kHz corner frequency appear in the jitter tolerance masks. The current SONET/SDH Option 2 specification are the result of an attempt to harmonize the SDH Option 1 and 2 specifications to the fullest extent possible; after this was done, the 12 kHz filter was retained only for jitter generation measurments for STM-16 (OC-48) and lower rates. This discussion is included here to point out that, while it is true that test equipment does exist that can make the jitter measurement with a 12 kHz high-pass filter, within the SONET/SDH specifications the filter is somewhat inconsistent with the jitter network limit and jitter tolerance specifications. It also has been pointed out to the editor that an RMS measurement might be easier to make than a peak-to-peak measurement. Here too, all the current SONET/SDH specifications are specified in terms of peak-to-peak jitter.

Comment 80 - 3

- Comment – cont.
  - The above material is background information. The specification should be set to a value that supports acceptable network performance at acceptable cost; in this case acceptable performance means that resulting jitter at the egress can be easily filtered by an endpoint filter. In addition, the test should be practical to make, preferably with existing test equipment.
  - Note that the comment is not directed against the wander specification in B.1.3.2. As indicated earlier, the dominant noise at lower frequencies, is FFM. The TDEV, ADEV, and PTPDEV masks in B.1.3.2 all reflect FFM, and are at the level of the noise model used in the earlier simulations.
- The following slides contain background and discussion
- Note that the reason for limiting jitter here is to enable the jitter requirements for applications to be met with reasonable endpoint filtering
  - We are not trying to ensure clock recovery with acceptable BER (other jitter specs do that)
  - We are not trying to prevent buffer overflow in chains of clocks (at the physical layer) or regenerators (we don't have these in 802.1AS)

Comment 80 - 4

- The clock noise model was taken from a model used in previous simulations in September and November, 2005
  - http://www.ieee802.org/3/re_study/public/200509/garner_1.pdf
  - http://www.ieee802.org/3/re_study/public/200511/20051114-garner-synch-simul.pdf
- Clock noise is modeled as a combination of 3 components:
  - White phase modulation (WPM, i.e., white phase noise)
  - Flicker phase modulation (FPM)
  - Flicker frequency modulation (FFM)

Comment 80 - 5

- Can characterize these noise processes by the respective one-sided power spectral density (PSD)

$$S_\varphi(f) = 4\int_0^\infty R_\varphi(\tau)\cos 2\pi f\tau\, d\tau$$

$$R_\varphi(\tau) = \int_0^\infty S_\varphi(f)\cos 2\pi f\tau\, df$$

$$R_\varphi(\tau) = E[\varphi(t)\varphi(t+\tau)]$$

$$R_\varphi(0) = E[\varphi^2(t)]$$

where

$\varphi(t)$ = phase noise process (units are usually rad, deg, or UI)

$S_\varphi(f)$ = power spectral density (units are usually rad^2/Hz or UI^2/Hz)

$R_\varphi(\tau)$ = autocorrelation function (units are usually rad^2 or UI^2)

$R_\varphi(0)$ = mean square (= variance for zero mean process, units are usually rad^2 or UI^2)

Comment 80 - 6

- For the above noise processes, the PSDs are
  - FFM: $S(f) = A/f^3$
  - FPM: $S(f) = B/f$
  - WPM: $S(f) = C$
- In the above, A, B, and C are constants.
- Note that for both WPM and FPM a noise bandwidth must be specified (because the integral over frequency diverges as $f \rightarrow \infty$
- The actual PSD was taken from the document: *Jitter and Signal Noise in Frequency Sources*, Raltron, Application Note, available at http://www.raltron.com/
  - In this document, PSD is specified using the related 2-sided spectrum $L(f)$
  - $S_\varphi(f) = 2L(f)$ (units of $L(f)$ are usually dBc/Hz)
- The spectrum is given on the following slide

Comment 80 - 7

Frequency (Hz), from carrier	$L(f)$, in dBc/Hz	$S_\varphi(f)$ in rad^2/Hz
10	-40	2×10^{-4}
100	-70	2×10^{-7}
1000	-100	2×10^{-10}
10000	-120	2×10^{-12}

To match this, we take:

$A = 0.2$ rad^{2}Hz^2 (i.e., the FFM term matches the PSD at 10 Hz)
$B = 2 \times 10^{-7}$ rad^2 (i.e., the FPM term matches the PSD at 1000 Hz
$C = 2 \times 10^{-12}$ rad^2/Hz (i.e., the WPM term matches the PSD at 10 kHz

The full PSD is plotted on the following slide

Comment 80 - 8

Example Clock Phase Noise Specification Provided in [7] (data in [7] does not extend above 10 kHz; PSD is assumed flat for higher frequencies with the 10 kHz value)

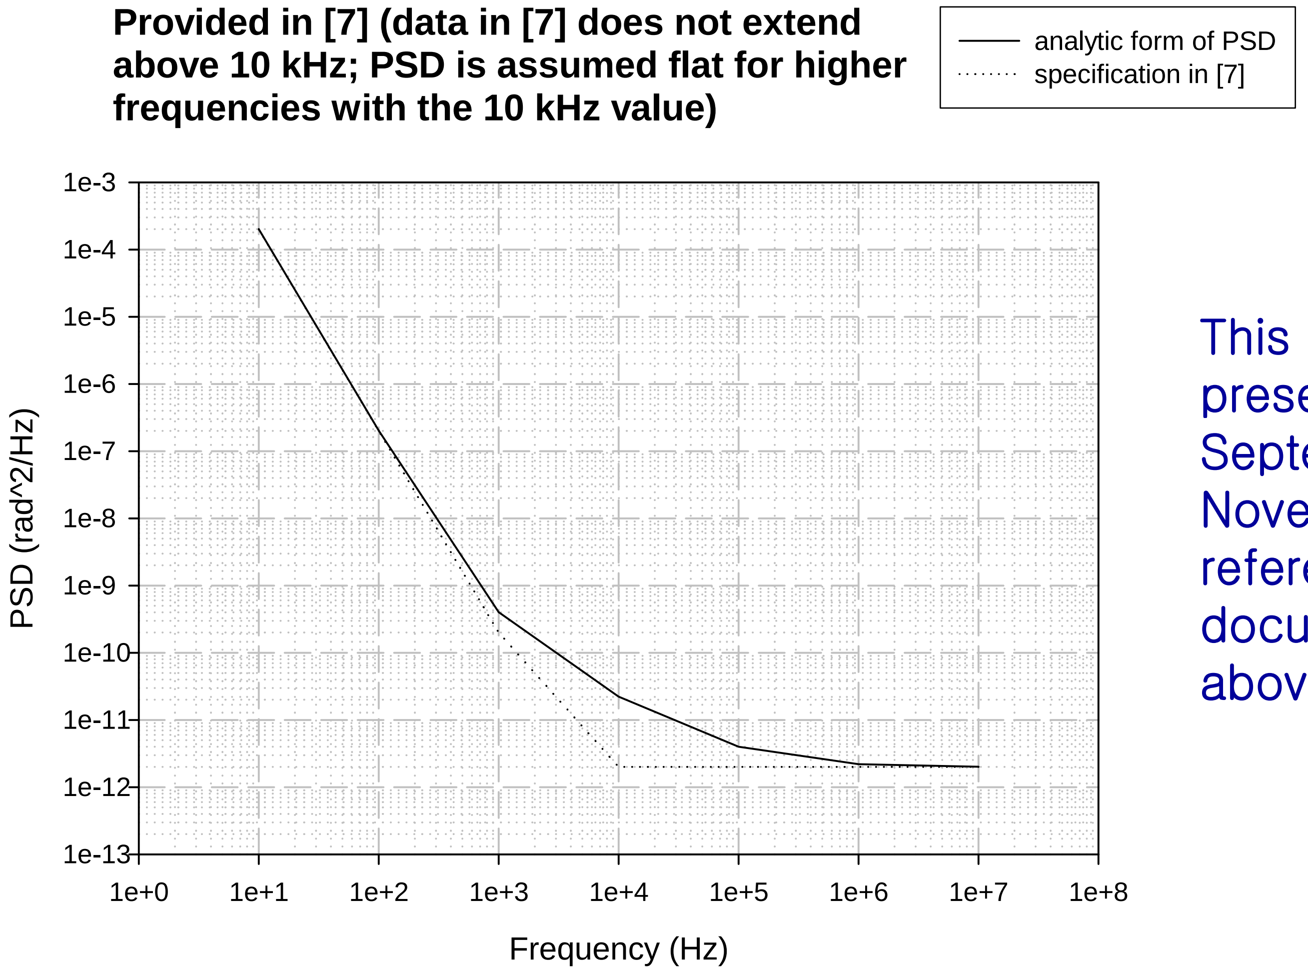

This plot is taken from presentations of September and November, 2005; reference [7] is the document cited above.

Comment 80 - 9

- Note that in this model, the FFM component dominates at low frequencies, i.e., in the wander region
  - This is typical for oscillators
  - The wander generation requirement of B.1.3.2 covers frequencies below 10 Hz
  - It is stated in terms of TDEV (with corresponding ADEV and PTPDEV shown as well)
    - These are given for observation intervals of 0.05 s and longer
    - If a 10 Hz anti-aliasing filter has been used to remove frequency components in the jitter region (i.e., > 10 Hz), then 10 Hz becomes the Nyquist rate, and the corresponding sampling interval is 0.05 s
  - For FFM (τ = observation interval)
    - TDEV is proportional to τ^1
    - ADEV is proportional to τ^0
    - PTPDEV is proportional to τ^1
- In the following slides, we focus on jitter, because that was the focus of this comment

Comment 80 - 10

- In B.1.3, jitter is measured through a 10 Hz high-pass measurement filter
- For a 25 MHz oscillator (ν_0 = 25 MHz), the Nyquist rate is 12.5 MHz
- The mean-square jitter is given by

$$\sigma^2 = \int_{10}^{\nu_0/2} S_\varphi(f)\,df$$

- The plot on slide 21 indicates that, between 10 Hz and 12.5 MHz, the FFM is the dominant noise component
- The comment indicates that the jitter was incorrectly computed using the WPM component
  - The resulting rms jitter was considerably less than 1 ns; this was rounded to 1 ns, which is still very small compared to the 40 ns phase measurement granularity
- In any case, the above is incorrect; we now compute the jitter due to each component

Comment 80 - 10

❑FFM component

$$\sigma^2 = \int_{10}^{\nu_0/2} \frac{A}{f^3} df = -\frac{A}{2f^2}\bigg|_{10}^{\nu_0/2} \cong \frac{A}{200\mathrm{Hz}^2} = 10^{-3}\,\mathrm{rad}^2 = \frac{10^{-3}}{(2\pi)^2}\mathrm{UI}^2 = 2.53\times10^{-5}\,\mathrm{UI}^2$$

$$\sigma(\mathrm{rms}) \cong 5\,\mathrm{mUI} = 0.2\ \mathrm{ns}$$

❑FPM component

$$\sigma^2 = \int_{10}^{\nu_0/2} \frac{B}{f} df = B\ln f\big|_{10}^{\nu_0/2} = B\ln\frac{\nu_0}{20} = B\ln(1.25\times10^6) = 2.81\times10^{-6}\,\mathrm{rad}^2 = \frac{2.81\times10^{-6}}{(2\pi)^2}\mathrm{UI}^2 = 7.11\times10^{-8}\,\mathrm{UI}^2$$

$$\sigma(\mathrm{rms}) \cong 0.27\,\mathrm{mUI} = 10.8\ \mathrm{ps}$$

❑WPM component

$$\sigma^2 = \int_{10}^{\nu_0/2} C\,df = C\left(\frac{\nu_0}{2} - 10\right) \cong 2.5\times10^{-5}\,\mathrm{rad}^2 = \frac{2.5\times10^{-5}}{(2\pi)^2}\mathrm{UI}^2 = 6.3\times10^{-7}\,\mathrm{UI}^2$$

$$\sigma(\mathrm{rms}) \cong 0.8\,\mathrm{mUI} = 32\ \mathrm{ps}$$

Comment 80 - 11

- The above indicate that the FFM component gives the largest contribution to jitter measured from 10 Hz.
- Next, the relation between rms and peak-to-peak jitter must be considered. The above model assumes phase noise has an underlying probability distribution
  - In this case, the peak-to-peak would be taken as some number of standard deviations (i.e., some multiple of the rms jitter), and the value of the multiplier along with the distribution determine the particular quantile that the peak-to-peak value corresponds to
  - For example, for a Gaussian distribution
    - 3σ corresponds to 0.9973 quantile, or 2.7×10^{-3} probability of exceeding the peak-to-peak value (in either the positive or negative direction)
    - 6σ corresponds to 1.97×10^{-9} probability of exceeding the peak-to-peak value (in either the positive or negative direction)
    - 9σ corresponds to 2.26×10^{-19} probability of exceeding the peak-to-peak value (in either the positive or negative direction)
    - 14σ corresponds to 1.56×10^{-44} probability of exceeding the peak-to-peak value (in either the positive or negative direction)

Comment 80 - 12

- The results show that, for the above model, the peak-to-peak jitter will certainly not exceed 3 ns (i.e., just over 14σ), which is still small compared to the 40 ns phase measurement granularity
- Note: the phase measurement granularity and jitter/wander are separate quantities
  - Phase measurement granularity refers to the 40 ns spacing of the clock edges
    - An event that we want to timestamp must be referred to the most recent edge, which means that the error can be as much as 40 ns
  - Jitter and wander refer to the deviation of the clock edges from their ideal positions in time
    - This gives rise to an additional error in a timestamp, because the most recent clock edge does not occur at its ideal time

Comment 80 - 13

- It also has been pointed out to the editor that the 10 Hz high-pass measurement filter may not be practical to test with, using existing test equipment
  - A value of 12 kHz has been suggested
  - Note that 12 kHz as a high-pass measurement filter is used to measure SONET (i.e., SDH Option 2) jitter generation; this is a long-standing measurement
    - Specified in both ANSI T1.105.03 and ITU-T Rec. G.783, for the rates
      - OC-1
      - OC-3 (STM-1)
      - OC-12 (STM-4)
      - OC-48 (STM-16)
      - **Not used** for STM-64 and STM-256
  - The 12 kHz filter is not used in SDH Option 1
    - Option 1 is for networks optimized for the 2048 kbit/s hierarchy
    - Option 2 is for networks optimized for the 1544 kbit/s hierarchy

Comment 80 - 14

- The 12 kHz highpass filter is not used for network interface (i.e., accumulated) jitter, nor is the value used in jitter tolerance masks (in either Option 1 or Option 2)
- Current SDH Option 2 specifications are the result of an attempt to harmonize the Option 1 and 2 specifications to the fullest extent possible; after this was done, the 12 kHz filter was retained only for jitter generation measurements, and only for the lower rates
- This discussion is included here to point out that, while it is true that test equipment does exist that can make the jitter measurement with a 12 kHz high-pass filter, within the SONET/SDH specifications the filer is somewhat inconsistent with the network limits and jitter tolerance specifications.
- It also has been pointed out to the editor that an rms measurement might be easier to make than a peak-to-peak measurement
  - Here too, all the current SONET/SDH specificaitons are in terms of peak-to-peak jitter measurements

Comment 80 - 15

- In summary, the following items must be resolved
  - What should the high-pass jitter measurement filter bandwidth be?
  - Should the spec be on peak-to-peak or rms jitter?
  - What should the jitter limit be?
- The answers to these questions will depend on
  - What do we assume for the magnitude and type of phase noise?
  - What is practical for existing test equipment?
- It is likely desirable to not have the jitter measurement filter corner frequency too high, as the higher-frequency jitter is easily filtered by endpoint filters and won't impact the jitter seen by the application timing appreciably
  - This means we should consider the MTIE masks for the various applications
- They are reproduced on the following slides

Comment 80 - 16

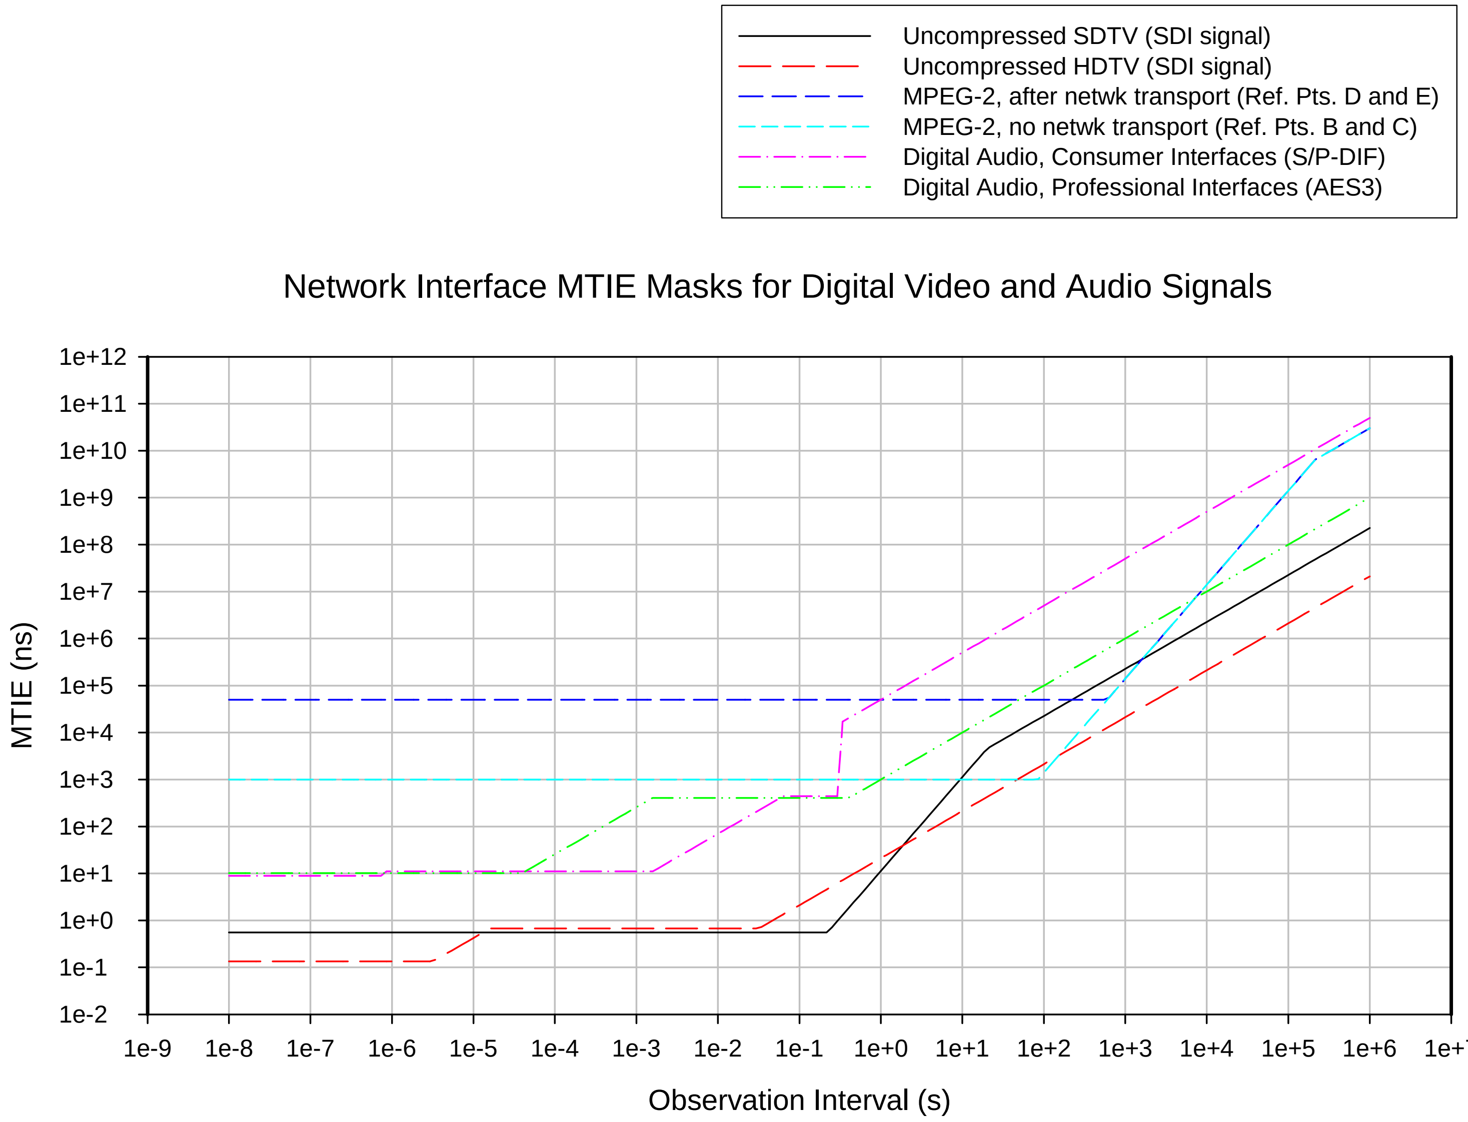

Comment 80 - 17

Requirement	Uncompressed SDTV	Uncompressed HDTV	MPEG-2, with network transport	MPEG-2, no network transport	Digital audio, consumer interface	Digital audio, professional interface
Wide-band jitter (UIpp)	0.2	1.0	50 μs peak-to-peak phase variation requirement (no measurement filter specified)	1000 ns peak-to-peak phase variation requirement (no measurement filter specified)	0.25	0.25
Wide-band jitter meas filt (Hz)	10	10			200	8000
High-band jitter (UIpp)	0.2	0.2			0.2	No requirement
High-band jitter meas filt (kHz)	1	100			400 (approx)	No requirement
Frequency offset (ppm)	±2.79365 (NTSC) ±0.225549 (PAL)	±10	±30	±30	±50 (Level 1) ±1000 (Level 2)	±1 (Grade 1) ±10 (Grade 2)
Frequency drift rate (ppm/s)	0.027937 (NTSC) 0.0225549 (PAL)	No requirement	0.000278	0.000278	No requirement	No requirement

Comment 80 - 18

- The uncompressed video masks have a corner frequency of 10 Hz
- The uncompressed SDTV mask has the equivalent of a corner frequency at 1 – 2 Hz due to the frequency drift requirement
- Both masks have peak-to-peak limits that are less than 1 ns
  - However, we have sent a liaison to SMPTE asking about the applicability of these requirements
- The audio masks have corner frequencies of 200 Hz (consumer) and 8 kHz (professional), with peak-to-peak limits around 10 ns
- Note that if we know the noise type, or make a conservative assumption, then we can choose a measurement filter with corner frequency greater than 10 Hz but still constrain the jitter from 10 Hz to 12.5 MHz